Reconfigurable Computing: Hardware Accelerators for Scientific Computing Applications

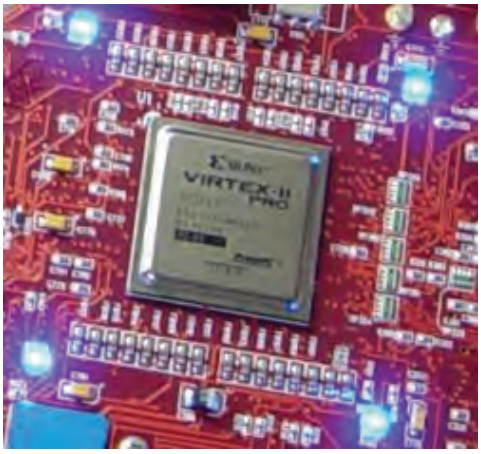

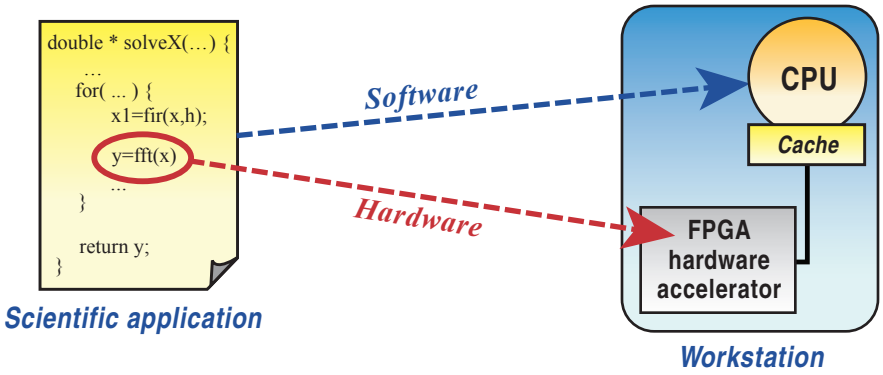

Figure 1. In reconfigurable computing, complex operations are ported to hardware and processed by field-programmable gate arrays (FPGAs) to improve performance.

To meet the scientific community's constant demand for increased computing power, researchers at Sandia National Laboratories are actively investigating emerging technologies that can be leveraged to improve the capabilities of high-performance computing (HPC) systems. While the traditional approach to enhancing performance has been to scale the number of compute nodes in an HPC system, an alternative method is to add hardware accelerators to individual nodes in a system.

These accelerators can perform certain scientific computing operations much more efficiently than general-purpose CPU and therefore represent opportunities to boost node performance. Multiple commercial products can be used as hardware accelerators, including the Cell (IBM/Sony/Toshiba), XMONARCH (Raytheon), and CSX (ClearSpeed) architectures.

However, these technologies are still emerging and lack widespread support. A more attractive technology for accelerator research that is available today is reconfigurable hardware.

Reconfigurable hardware devices such as field-programmable gate arrays (FPGAs) can be programmed to emulate custom digital hardware circuitry. Current generation FPGAs house up to 10 million logic gates in a single chip and can be reprogrammed with new circuitry in a few milliseconds. With this in mind, reconfigurable computing (RC) researchers are using FPGAs as an affordable way to implement an algorithm in hardware instead of software and have achieved significant application speedups. In support of this work, HPC vendors are beginning to include FPGA accelerators in their system architectures.

SAND2006-2517P
April 2006

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Sandia National Laboratories is currently involved in RC research that seeks to make FPGAs a practical option for accelerating scientific applications. In this work, Sandia scientists are attacking three formidable research problems of great interest to the RC community: providing support for floating-point arithmetic, automating the adaptation of algorithms to hardware, and integrating FPGAs into HPC systems.

Floating-Point Arithmetic. Current-generation FPGAs lack native support for floating-point operations. Given the prevalence of floating-point calculations in scientific applications, Sandia researchers have constructed a series of floating-point units that are compliant with the IEEE-754 standard for binary floating-point arithmetic. These units, which can be used in the Xilinx Virtex II/Pro family of FPGAs, operate at speeds as high as 200 MHz and are deeply pipelined to maximize performance. The units are compact, allowing up to 70 independent floating-point operations to be placed in a single FPGA.

Adapting Algorithms to Hardware. To combat the fact that hardware is significantly more time-consuming to develop than software, Sandia researchers have constructed tools to help automate the process of translating a software algorithm to hardware. These tools analyze the structure of an algorithm and extract a data-flow graphical representation of the computation. This graph is manipulated to find the optimal ordering of operations for a given FPGA technology. The tools then generate a synthesizable hardware netlist that can be compiled and programmed into an FPGA.

Integrating FPGAs into HPC Systems. In order for FPGAs to be leveraged as accelerators, they must first be integrated into an HPC system. Sandia researchers are investigating two different strategies for system integration.

First, in a networked approach, FPGAs are connected to an HPC system through an existing communication network. Sandia researchers have developed a Gigabit Ethernet network interface for FPGAs that contains a built-in TCP offload engine. The advantage of this approach is that FPGA accelerators can be added to existing HPC systems with minimal modifications. A second approach to system integration is to use HPC systems that feature FPGA accelerators in the native architectures. For example, the Cray XD1 provides an architecture where each pair of AMD Opterons in the system is equipped with an FPGA coprocessor. Sandia researchers have constructed FPGA hardware and host software to enable applications to rapidly exchange data with the FPGAs.

Sandia's RC research effort has demonstrated that FPGAs can be used as hardware accelerators for scientific-computing applications. Early results have produced speedups that are 2 to 10 times faster than approaches that rely solely on software implementations of algorithms. We expect performance will continue to improve as the tools and hardware developed in this research effort are refined and as additional applications are adapted to hardware. This work, combined with research into other emerging accelerators, will help propel the scientific and engineering communities at Sandia National Laboratories as they continue to answer the challenging questions facing our nation.

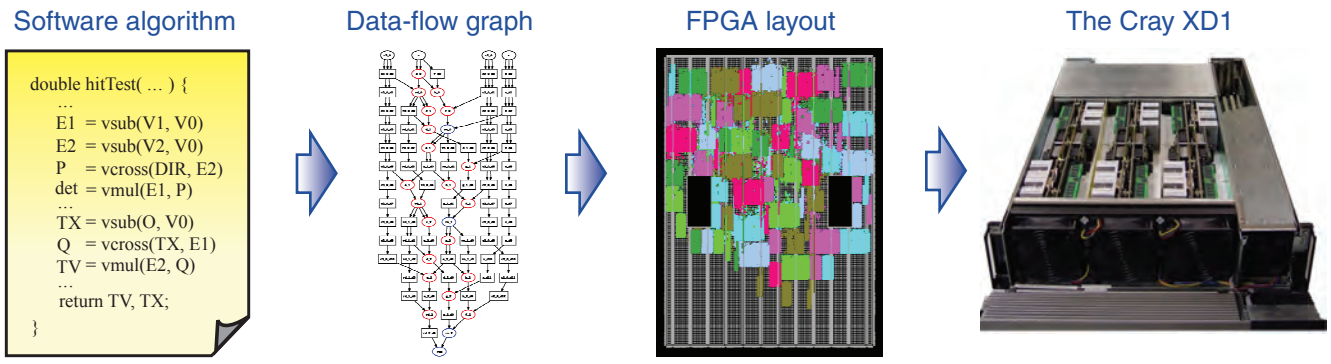

Figure 2. In reconfigurable computing, an algorithm is ported to hardware and then programmed into the FPGAs of a system such as the Cray XD1.

For more information, contact
Craig Ulmer (925) 294-6061